EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	0	"707".clas. and virtual adj processor same node same map same subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2006/09/22 10:38
S2	0	"707".clas. and virtual adj processor same node same map$4 same subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:06
S3	0	"709".clas. and virtual adj processor same node same map$4 same subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:07
S4	0	"703".clas. and virtual adj processor same node same map$4 same subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:06
S5	3	"703".clas. and virtual adj processor same node	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:06
S6	19	"709".clas. and virtual adj processor same node	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:09
S7	0	S6 and map$4 same subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:08
S8	2	S6 and map$4 and subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:10
S9	3	"710".clas. and virtual adj processor same node	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:09

EAST Search History

S10	16	"711".clas. and virtual adj processor same node	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:12
S11	5	S10 and map$4 and subset	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:10
S12	2	"713".clas. and virtual adj processor same node	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/09/21 18:12

Terms used **virtual processor tranaction** and **node** and **subjet** Found **30,034** of **185,178**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window
Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Update propagation in Bakunin data networks
Boris Kogan, Hector Garcia-Molina
December 1987 **Proceedings of the sixth annual ACM Symposium on Principles of distributed computing PODC '87**
Publisher: ACM Press
Full text available: pdf(1.47 MB) Additional Information: full citation, references, citings, index terms

2 Virtual nodes and units of distribution for distributed Ada
R. A. Volz
May 1989 **ACM SIGAda Ada Letters , Proceedings of the third international workshop on Real-time Ada issues IRTAW '89**, Volume X Issue 4
Publisher: ACM Press
Full text available: pdf(1.01 MB) Additional Information: full citation, references, citings, index terms

3 Location-aware mechanisms: Autonomous virtual mobile nodes
Shlomi Dolev, Seth Gilbert, Elad Schiller, Alex A. Shvartsman, Jennifer Welch
September 2005 **Proceedings of the 2005 joint workshop on Foundations of mobile computing DIALM-POMC '05**
Publisher: ACM Press
Full text available: pdf(178.37 KB) Additional Information: full citation, abstract, references, index terms

This paper presents a new abstraction for virtual infrastructure in mobile ad hoc networks. An Autonomous Virtual Mobile Node (AVMN) is a robust and reliable entity that is designed to cope with the inherent difficulties caused by processors arriving, leaving, and moving according to their own agendas, as well as with failures and energy limitations. There are many types of applications that may make use of the AVMN infrastructure: tracking, supporting mobile users, or searching for energy sourc ...

Keywords: ad hoc networks, distributed algorithms, fault-tolerance, location-aware, mobile networks, virtual infrastructure

4 Devirtualizable virtual machines enabling general, single-node, online maintenance
David E. Lowell, Yasushi Saito, Eileen J. Samberg

October 2004 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 32 , 38 , 39 Issue 5 , 5 , 11
Publisher: ACM Press

Full text available: pdf(174.01 KB) Additional Information: full citation, abstract, references, citings, index terms

Maintenance is the dominant source of downtime at high availability sites. Unfortunately, the dominant mechanism for reducing this downtime, cluster rolling upgrade, has two shortcomings that have prevented its broad acceptance. First, cluster-style maintenance over many nodes is typically performed a few nodes at a time, mak-ing maintenance slow and often impractical. Second, cluster-style maintenance does not work on single-node systems, despite the fact that their unavailability during mainte ...

Keywords: availability, online maintenance, planned downtime, virtual machines

5 On virtual memories and micronetworks

G. Jack Lipovski
March 1977 **ACM SIGARCH Computer Architecture News , Proceedings of the 4th annual symposium on Computer architecture ISCA '77**, Volume 5 Issue 7
Publisher: ACM Press
Full text available: pdf(891.29 KB) Additional Information: full citation, abstract, references, index terms

We propose to use the microcomputer in a network to share I/O resources such as printers and archival memories. A model of a network is developed where computers correspond to edges of a graph. This model reflects the desired characteristics of the microcomputer organization. The advantage of virtual memory in these microcomputers is discussed. Herein, pages in the virtual memory are packets in the network. Packets and requests for packets are generated by page faults and packets are stored ...

6 Cellular Disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum
December 1999 **ACM SIGOPS Operating Systems Review , Proceedings of the seventeenth ACM symposium on Operating systems principles SOSP '99**, Volume 33 Issue 5
Publisher: ACM Press

Full text available: pdf(1.93 MB) Additional Information: full citation, abstract, references, citings, index terms

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that leverages the existing operating system technology.In this paper we present a syste ...

7 Accelerating shared virtual memory via general-purpose network interface support

Angelos Bilas, Dongming Jiang, Jaswinder Pal Singh
February 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 1
Publisher: ACM Press

Full text available: pdf(178.88 KB) Additional Information: full citation, abstract, references, index terms, review

Clusters of symmetric multiprocessors (SMPs) are important platforms for high-performance computing. With the success of hardware cache-coherent distributed shared

memory (DSM), a lot of effort has also been made to support the coherent shared-address-space programming model in software on clusters. Much research has been done in fast communication on clusters and in protocols for supporting software shared memory across them. However, the performance of software virtual memory (SVM) is sti ...

Keywords: applications, clusters, shared virtual memory, system area networks

8 The higher radix hypercube as an interconnection and virtual network

T.-C. Lin, P. Gupta

February 1989 **Proceedings of the 17th conference on ACM Annual Computer Science Conference**

Publisher: ACM Press

Full text available: pdf(764.16 KB) Additional Information: full citation, abstract, references, index terms

The HIGHER RADIX HYPERCUBE (HRH) is investigated in this paper as an interconnection network for multiprocessors as well as a virtual network. The HRH is based on a radix higher than two, the radix for the hypercube. We have compared several topological parameters' for the HRH with the hypercube. It can be seen that HRH provides smaller diameter for a small price. A mapping and a partitioning algorithms are proposed for mapping different topologies to the HRH. In our algorithm, we first map ...

9 Cellular disco: resource management using virtual clusters on shared-memory multiprocessors

Kinshuk Govil, Dan Teodosiu, Yongqiang Huang, Mendel Rosenblum

August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available: pdf(287.05 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Despite the fact that large-scale shared-memory multiprocessors have been commercially available for several years, system software that fully utilizes all their features is still not available, mostly due to the complexity and cost of making the required changes to the operating system. A recently proposed approach, called Disco, substantially reduces this development cost by using a virtual machine monitor that laverages the existing operating system technology. In this paper we present a ...

Keywords: fault containment, resource managment, scalable multiprocessors, virtual machines

10 A comparative analysis of virtual versus physical process-migration strategies for distributed modeling and simulation of mobile computing networks

Kwun Han, Sumit Ghosh

August 1998 **Wireless Networks**, Volume 4 Issue 5

Publisher: Kluwer Academic Publishers

Full text available: pdf(252.81 KB) Additional Information: full citation, abstract, references, index terms

Improvements in processor power and diminishing processor costs coupled with the potential of asynchronous, distributed algorithms promise to expand the frontier of mobile computing networks. In general, a mobile computing network consists of semi-autonomous or autonomous stationary and mobile agents that perform local computations, cooperate, and communicate among themselves to achieve a desired objective. While the stationary entities are connected through a static interconnection network ...

11 The Wisconsin Wind Tunnel: virtual prototyping of parallel computers

Steven K. Reinhardt, Mark D. Hill, James R. Larus, Alvin R. Lebeck, James C. Lewis, David A. Wood

June 1993 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '93**, Volume 21 Issue 1

Publisher: ACM Press

Full text available: pdf(1.40 MB) Additional Information: full citation, abstract, references, citings, index terms

We have developed a new technique for evaluating cache coherent, shared-memory computers. The Wisconsin Wind Tunnel (WWT) runs a parallel shared-memory program on a parallel computer (CM-5) and uses execution-driven, distributed, discrete-event simulation to accurately calculate program execution time. WWT is a virtual prototype that exploits similarities between the system under design (the target) and an existing evaluation platform (the host). The host directly executes all target program ins ...

12 Using network interface support to avoid asynchronous protocol processing in shared virtual memory systems

Angelos Bilas, Cheng Liao, Jaswinder Pal Singh

May 1999 **ACM SIGARCH Computer Architecture News , Proceedings of the 26th annual international symposium on Computer architecture ISCA '99**, Volume 27 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(440.73 KB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

The performance of page-based software shared virtual memory (SVM) is still far from that achieved on hardware-coherent distributed shared memory (DSM) systems. The interrupt cost for asynchronous protocol processing has been found to be a key source of performance loss and complexity.This paper shows that by providing simple and general support for asynchronous message handling in a commodity network interface (NI), and by altering SVM protocols appropriately, protocol activity can be decoupled ...

13 Memory coherence in shared virtual memory systems

Kai Li, Paul Hudak

November 1986 **Proceedings of the fifth annual ACM symposium on Principles of distributed computing PODC '86**

Publisher: ACM Press

Full text available: pdf(773.45 KB) Additional Information: full citation, references, citings, index terms

14 A generalized processor sharing approach to flow control in integrated services networks: the single-node case

Abhay K. Parekh, Robert G. Gallager

June 1993 **IEEE/ACM Transactions on Networking (TON)**, Volume 1 Issue 3

Publisher: IEEE Press

Full text available: pdf(1.61 MB) Additional Information: full citation, references, citings, index terms, review

15 SoftFLASH: analyzing the performance of clustered distributed virtual shared memory

Andrew Erlichson, Neal Nuckolls, Greg Chesson, John Hennessy

September 1996 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the seventh international conference on Architectural support for programming languages and operating systems ASPLOS-VII**, Volume 31 , 30 Issue 9 , 5

Publisher: ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, citings, index terms

One potentially attractive way to build large-scale shared-memory machines is to use small-scale to medium-scale shared-memory machines as clusters that are interconnected with an off-the-shelf network. To create a shared-memory programming environment across the clusters, it is possible to use a virtual shared-memory software layer. Because of the low latency and high bandwidth of the interconnect available within each cluster, there are clear advantages in making the clusters as large as possi ...

16 High performance communications in processor networks

C. R. Jesshope, P. R. Miller, J. T. Yantchev

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture ISCA '89**, Volume 17 Issue 3

Publisher: ACM Press

Full text available: pdf(1.11 MB) Additional Information: full citation, abstract, references, citings, index terms

In order to provide an arbitrary and fully dynamic connectivity in a network of processors, transport mechanisms must be implemented, which provide the propagation of data from processor to processor, based on addresses contained within a packet of data. Such data transport mechanisms must satisfy a number of requirements - deadlock and livelock freedom, good hot-spot performance, high throughput and low latency. This paper proposes a solution to these problems, which allows deadlock free, ...

17 Helper threads via virtual multithreading on an experimental itanium® 2 processor-based platform

Perry H. Wang, Jamison D. Collins, Hong Wang, Dongkeun Kim, Bill Greene, Kai-Ming Chan, Aamir B. Yunus, Terry Sych, Stephen F. Moore, John P. Shen

October 2004 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , ACM SIGARCH Computer Architecture News , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 39 , 38 , 32 Issue 11 , 5 , 5

Publisher: ACM Press

Full text available: pdf(225.47 KB) Additional Information: full citation, abstract, references, citings, index terms

Helper threading is a technology to accelerate a program by exploiting a processor's multithreading capability to run ``assist'' threads. Previous experiments on hyper-threaded processors have demonstrated significant speedups by using helper threads to prefetch hard-to-predict delinquent data accesses. In order to apply this technique to processors that do not have built-in hardware support for multithreading, we introduce virtual multithreading (VMT), a novel form of switch-on-event user-level ...

Keywords: DB2 database, PAL, cache miss prefetching, helper thread, itanium processor, multithreading, switch-on-event

18 Bandwidth scheduling for wide-area ATM networks using virtual finishing times

Anthony Hung, George Kesidis

February 1996 **IEEE/ACM Transactions on Networking (TON)**, Volume 4 Issue 1

Publisher: IEEE Press
Full text available: pdf(777.18 KB) Additional Information: full citation, references, citings, index terms

19 Memory coherence in shared virtual memory systems

Kai Li, Paul Hudak
November 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 4
Publisher: ACM Press

Full text available: pdf(2.71 MB) Additional Information: full citation, abstract, references, citings, index terms, review

The memory coherence problem in designing and implementing a shared virtual memory on loosely coupled multiprocessors is studied in depth. Two classes of algorithms, centralized and distributed, for solving the problem are presented. A prototype shared virtual memory on an Apollo ring based on these algorithms has been implemented. Both theoretical and practical results show that the memory coherence problem can indeed be solved efficiently on a loosely coupled multiprocessor.

20 Virtual memory mapped network interface for the SHRIMP multicomputer

M. A. Blumrich, K. Li, R. Alpert, C. Dubnicki, E. W. Felten, J. Sandberg
April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture ISCA '94**, Volume 22 Issue 2
Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references, citings, index terms

The network interfaces of existing multicomputers require a significant amount of software overhead to provide protection and to implement message passing protocols. This paper describes the design of a low-latency, high-bandwidth, virtual memory-mapped network interface for the SHRIMP multicomputer project at Princeton University. Without sacrificing protection, the network interface achieves low latency by using virtual memory mapping and write-latency hiding techniques, and obtains high bandw ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Results for "((virtual processor and node)<in>metadata)" e-mail
Your search matched **6** of **1415139** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

View Session History

New Search

Modify Search

((virtual processor and node)<in>metadata) Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Key

IEEE JNL	IEEE Journal or Magazine
IEE JNL	IEE Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IEE CNF	IEE Conference Proceeding
IEEE STD	IEEE Standard

view selected items Select All Deselect All

1. **An enhanced parallel toroidal lattice architecture for large scale neural ne**
Fujimoto, Y.; Fukuda, N.;
Neural Networks, 1989. IJCNN., International Joint Conference on
18-22 June 1989 Page(s):614 vol.2
Digital Object Identifier 10.1109/IJCNN.1989.118454
AbstractPlus | Full Text: PDF(80 KB) IEEE CNF
Rights and Permissions

2. **Shared-memory architecture to implement a high-connectivity processin**
Ancona, F.; Anguita, D.; Rovetta, S.; Zunino, R.;
Electronics Letters
Volume 31, Issue 22, 26 Oct. 1995 Page(s):1903 - 1904
AbstractPlus | Full Text: PDF(212 KB) IEE JNL

3. **An Infrastructure for Efficient Parallel Job Execution in Terascale Compu Environments**
Moreira, J.E.; Waiman Chan; Fong, L.L.; Franke, H.; Jette, M.A.;
Supercomputing, 1998. SC98. IEEE/ACM Conference on
07-13 Nov. 1998 Page(s):50 - 50
Digital Object Identifier 10.1109/SC.1998.10026
AbstractPlus | Full Text: PDF(416 KB) IEEE CNF
Rights and Permissions

4. **Three-dimensional modeling of electrical defibrillation on a massively pa**
Hutchinson, S.A.; Gao, S.; Ai, L.; Ng, K.T.; Deale, O.C.; Cahill, P.T.; Lerman, E
Computers in Cardiology 1992. Proceedings.
11-14 Oct. 1992 Page(s):343 - 346
Digital Object Identifier 10.1109/CIC.1992.269361
AbstractPlus | Full Text: PDF(304 KB) IEEE CNF
Rights and Permissions

5. **Implementing halt on failure processors**
Macdonald, R.; Shoja, G.;
Communications, Computers and Signal Processing, 1993., IEEE Pacific Rim
Volume 1, 19-21 May 1993 Page(s):272 - 275 vol.1
Digital Object Identifier 10.1109/PACRIM.1993.407171
AbstractPlus | Full Text: PDF(324 KB) IEEE CNF

Rights and Permissions

☐ **6. Three-dimensional finite-difference bidomain modeling of homogeneous on a data-parallel computer**
Saleheen, H.I.; Claessen, P.D.; Ng, K.T.;
Biomedical Engineering, IEEE Transactions on
Volume 44, Issue 2, Feb. 1997 Page(s):200 - 204
Digital Object Identifier 10.1109/10.552249

AbstractPlus | References | Full Text: PDF(136 KB) **IEEE JNL**
Rights and Permissions